

The Budgeted Transportation Problem

S. Kapoor^{*} and M. Sarwat[†]

Abstract

Consider a transportation problem with sets of sources and sinks. There are profits and prices on the edges. The goal is to maximize the profit while meeting the following constraints; the total flow going out of a source must not exceed its capacity and the total price of the incoming flow on a sink must not exceed its budget. This problem is closely related to the generalized flow problem.

We propose an auction based primal dual approximation algorithm to solve the problem. The complexity is $O(\epsilon^{-1}(n^2 + n \log m)m \log U)$ where n is the number of sources, m is the number of sinks, U is the ratio of the maximum profit/price to the minimum profit/price.

We also show how to generalize the scheme to solve a more general version of the problem, where there are edge capacities and/or the profit function is concave and piecewise linear. The complexity of the algorithm depends on the number of linear segments, termed \mathcal{L} , of the profit function.

1 Introduction

The Transportation Problem is a fundamental problem in Computer Science. It was initially formulated to represent the problem of transporting goods from warehouses to customers. The mincost version of the problem is also known as the Hitchcock problem after one of its early formulators [15]. It is well-known that while the Transportation problem can be seen as a special case of the mincost flow problem, the latter can also be transformed into the former [21].

The first primal-dual algorithm, called the Hungarian method [17], was also proposed for the Assignment Problem which happens to be a special case of the Transportation Problem. More recently, the Transportation Problem was used by Bertsekas [3] and Goldberg and Tarjan [11] as a natural framework for applying auction based algorithms.

The typical transportation model assumes that there is no loss of goods as the goods are transported from the sources to the sinks. In real life, there is often a loss involved. For example, in case of electricity there is loss due to resistance and in case of oil, due to evaporation. In case of a delicate merchandise, there might breakage. Sometimes, there may even be a gain (e.g. transfer of money: currency conversion). The loss or gain can be modeled as a price function on the transport from the source to the sink.

As an example, consider a set of depots and a set of retail stations. The depots have a certain supply of goods, while the retail stations have a limit on their intake capacity. For every unit transported between a pair of depot and station, a certain amount of profit is made which depends on selling price, cost of transportation etc. In addition, consider the situation where goods transported suffer a loss proportional to the number of units transported. This factor is unique to each pair as it depends on the mode of transfer and distance etc. In order to measure

^{*}Computer Science, IIT, Chicago, IL, e-mail: kapoor@iit.edu

[†]Amazon.com, Work done at Computer Science, IIT, Chicago, IL, e-mail sarwat@iit.edu

the net incoming amount at the stations, we have to account for this loss. The intake capacity limit is therefore, on a weighted sum rather than a simple sum.

In a more recent context, keyword based advertising is widely applied by internet based service companies like Google and Yahoo. Consider the following model for assigning advertisement space to the bidders: each bidder, say j , modeled by a sink, specifies its total budget b_j and the price p_{ij} it is willing to pay for the i th keyword, specified by source i . An estimate of the number of times a keyword would be invoked in web searches provides a capacity a_i for the i th keyword. Assignment of a keyword i to a bidder j provides a profit c_{ij} . The goal is to maximize the revenue of the service provider without exceeding the bidders budget. A version of the problem restricted to an integral setting may be found in [2]. The online version of the problem has been considered in [19].

In order to model the above and similar problems, we propose a generalization of the transportation problem: the Budgeted Transportation Problem (BTP) which is closely related to Generalized Flow Problem. We then present an auction based approximation algorithm. We also show how to extend the technique to further generalizations of the problem.

Like the conventional Transportation Problem, the problem is modeled on a weighted bipartite graph. In our version of the problem we consider maximizing the profit in transporting/assigning goods from the sources to the sinks, subject to the following conditions: Each source, i has a *supply* of goods, bounded by a capacity function a_i , and each sink j has a bound on the incoming set of goods, specified by a *budget* b_j . The difference from the conventional transportation problem is that the budget is a bound on the weighted sum of the incoming goods as follows: each source-sink pair (i, j) has a *price* function p_{ij} and the bound b_j is on the sum $\sum_i p_{ij} f_{ij}$ where f_{ij} is the amount of goods transported/assigned from source i to sink j . The goal is to maximize the sum of profits of transporting/assigning goods from sources to sinks, given that the profit obtained by transporting/assigning f_{ij} units is $c_{ij} f_{ij}$. The problem can be formally stated as follows:

Given a bi-partite graph $G(S, T)$, capacity $a : S \rightarrow \mathbb{Z}^+$, budget $b : T \rightarrow \mathbb{Z}^+$, price $p : S \times T \rightarrow \mathbb{Z}^+$, profit $c : S \times T \rightarrow \mathbb{Z}^+$ the budgeted transportation problem is to find a flow function $f : S \times T \rightarrow \mathbb{R}^+$ so as to maximize $\sum_{i \in S, j \in T} c_{ij} f_{ij}$ subject to capacity constraints, $\sum_{j \in T} f_{ij} \leq a_i$ and budget constraints $\sum_{i \in S} p_{ij} f_{ij} \leq b_j, \forall j \in T$. We let $|S| = n$ and $|T| = m$.

Note that, for simplicity, the price function, like the other functions, is assumed to be integral in the paper. This could be replaced by a price function to the set of rationals.

We describe an approximation scheme for the budgeted transportation problem. Our algorithm is a primal-dual auction algorithm. Auction algorithms have been utilized before in the solution of combinatorial problems [3, 11] including the classical transportation problem [5], and more recently for finding market equilibrium [9, 10]. Bertsekas et al [4] apply the auction mechanism to solve the generalized flow problem but the complexity of that algorithm is not polynomial.

While the auction technique has been used before, the application in the current context is different. The auction mechanism is used to realize a set of tight edges, i.e. edges which satisfy complementary slackness conditions. Paths and cycle are found in a subgraph in these set of tight edges. The flow path/cycle could either be flow generating or create a flow reduction. The primal-dual nature of our approach allows us to push flow along these paths or cycle without worrying about the profit of the path/cycle. The cycle can be eliminated in linear time. Our algorithm may be interpreted to have achieved a version of cycle-canceling [25] without having to compute the associated profit.

We achieve a complexity of $O(\epsilon^{-1}(n^2 + n \log m)m \log U \log m)$, where n is the number of

sources, m is the number of sinks and U is $\frac{\max_{ij}(\frac{p_{ij}}{c_{ij}})}{\epsilon \min_{ij}(\frac{p_{ij}}{c_{ij}})}$. The problem can also be modeled as a generalized flow problem. This can be done by adding a super source connected to all the sources and a supersink connected to all the sinks and minimizing the negative of the profit function. By using the technique of scaling as in [18], the dependancy on $\log U$ can be replaced by $\log(nm/\epsilon)$. We will thus choose to ignore terms involving $\log U$, and in fact $\log n$, when comparing our results with other results.

There are numerous results for maximum/mincost generalized flow and related problems [24, 7, 25, 23, 20, 12, 13]. Some of the currently best known FPTAS for the generalized flow problems are of complexity $\tilde{O}(\log \epsilon^{-1} E(E + V \log I))$ [7, 24, 23] for maximum generalized flow and $\tilde{O}(\epsilon^{-2} E^2 V J)$ [7] and $\tilde{O}(\log \epsilon^{-1} E^2 V^2)$ [25] for minimum cost generalized flow. Here, $I = \log M$, $J = \log \log M + \log \epsilon^{-1}$ and M is the largest integer in cost representation. For the special case, when there are no flow generating cycles, the complexities are $\tilde{O}(\epsilon^{-2} E^2)$ and $\tilde{O}(\epsilon^{-2} E^2 J)$ for maximum and minimum cost generalized flows respectively [7], where V is the number of nodes and E is the number of edges in the given graph.

A straight-forward application of the best known generalized flow approximation algorithm to this problem would require $\tilde{O}(\log \epsilon^{-1} (mn)^2 (m+n)^2)$ time using the minimum cost generalized flow algorithm in [25] or $\tilde{O}(\epsilon^{-2} (mn)^2 (m+n) J)$ time using the algorithm in [7]. Our algorithm has a better complexity than the first when $\epsilon > 1/m^3$ and is better than the second for all ranges of ϵ .

If we use the packing algorithm from [8] than one can solve this problem in $\tilde{O}(\epsilon^{-2} (n + m)nm)$ since we have $m + n$ rows and mn columns in the primal constraint matrix. Other known combinatorial algorithms for packing [22, 14] are also all dependent on $1/\epsilon^2$. After the completion of this work we became aware of another result [16] which provides a randomized $(1 + \epsilon)$ approximate solution to the packing problem in time $O(N + (r + c) \log(N)/\epsilon^2)$ where N is the number of non-zero entries in the specification of the packing program and $r + c$ is the number of rows and columns in the matrix specifying the packing program. As applied to the budgeted transportation problem this algorithm would require $O(nm(\log nm)/\epsilon^2)$ and would be faster than our approach when $\epsilon \geq 1/n$.

Comparing with algorithms that have a dependence on $1/\epsilon$, one recent algorithm for packing problems is by Bienstock and Iyengar [6], which reduces the packing problem to solve a sequence of quadratic programming which in turn are approximated by linear programs. Consequently, the complexity will have high polynomial dependance on n and m . We achieve a combinatorial algorithm with complexity dependant on $O(1/\epsilon)$ and a low polynomial in n and m .

We extend our algorithm to solve a generalization of BTP called the Budgeted Transshipment Problem (BTS). It is similar to BTP except for added edge capacity constraints. We achieve the same approximation with no additional time complexity, i.e $O(\epsilon^{-1} (n^2 + n \log m) m \log U \log m)$. This is of improved efficiency, since, again, to compare this with existing results: the result in [25] would imply a complexity of $\tilde{O}(\log \epsilon^{-1} (mn)^2 (m+n)^2)$, [7] would imply a complexity of $\tilde{O}(\epsilon^{-2} (mn)^2 (m+n) J)$ while [8] would imply a complexity of $\tilde{O}(\epsilon^{-2} (n + m + mn) nm)$ as the number of constraints is now $m + n + mn$. The results from [16] would, again, be of complexity $O(nm(\log nm)/\epsilon^2)$.

In section 2.1 we present a basic auction based algorithm and show its correctness. In section 2.2 we describe a modified algorithm and discuss its convergence. Since the algorithm is similar to the algorithm for the generalized problem, BTS, we demonstrate an outline of the algorithm. Section 4 extends the result to BTS. In section 4.1 we present a detailed algorithm for BTS (which also applies to BTP) and in subsequent sections the proofs of correctness and complexity of the algorithm. The piece-wise linear case is discussed in section 5. We conclude in section 6. We also show a relationship between BTP and minimum cost generalized flow problems in the

appendix (section 7).

2 Auction Based Algorithm

2.1 The Basic Method

In this section, we describe a primal-dual auction framework to approximate the maximum profit budgeted transportation problem (BTP) to within a factor of $(1 - \epsilon)$. The algorithm is a locally greedy algorithm where the sources compete against each other for sinks by bidding with lower effective profit, a quantity defined using dual variables. This algorithm will serve to illustrate the methodology. In order to prove complexity bounds we will, in the next section, refine the algorithm further.

We first model the problem using an LP and also define its dual.

$$\begin{aligned}
& \text{maximize} && \sum_{i \in S, j \in T} c_{ij} f_{ij} \\
& \text{subject to :} && \\
& && \sum_{j \in T} f_{ij} \leq a_i \quad \forall i \in S \\
& && \sum_{i \in S} p_{ij} f_{ij} \leq b_j \quad \forall j \in T \\
& && f_{ij} \geq 0 \quad \forall i \in S, j \in T
\end{aligned}$$

In the remainder of the paper, the variables i and j will refer to a source i in S and sink j in T , respectively.

The dual of the above program is

$$\begin{aligned}
& \text{minimize} && \sum_{i \in S} \alpha_i a_i + \sum_{j \in T} \beta_j b_j \\
& \text{subject to :} && \\
& && \alpha_i \geq c_{ij} - p_{ij} \beta_j \quad \forall ij \\
& && \alpha_i, \beta_j \geq 0
\end{aligned}$$

The following variables based on primal and dual variables are used in the algorithm.

- α_i : the dual variable associated with each source.
- β_j : the dual variable corresponding to the budget constraint on the sink. An algorithmic interpretation of β_j is that it represents the value of sink j . It increases as different sources compete in order to ship flow to it. At every instance of the algorithm, flow may be pushed onto the sink j from a source at a given value of β_j . We term this an assignment of the source at valuation β_j . Further, at any instance of the algorithm sources are assigned to sinks at two prices (values), some at value β while others are assigned at a **companion valuation** β'_j where $\beta_j = \beta'_j(1 + \epsilon)$.
- y_{ij} : the **valuation** of sink j when assigned to source i to ship the flow from i to j . It equals β_j when source i ships flow to sink j . Subsequently, as β_j rises, it becomes equal to β'_j .

- s_i : the supply (or surplus) remaining at the source i . s_i is evaluated as $s_i = a_i - \sum_j f_{ij}$ and is implicitly updated in the algorithm.
- d_j : the amount remaining from the budget at sink j . d_j is evaluated as $d_j = b_j - \sum_i p_{ij} f_{ij}$.

Because the input is assumed to be integral, we note that all the variables are rationals, with a bound on the absolute value of the denominators.

The initial and primal dual feasible solutions are obtained by the initialization procedure 2.1 **Initialize**. The algorithm maintains a pair of feasible dual and primal solutions. It works to modify them until the following complementary slackness conditions, which define optimality, are exactly or approximately satisfied.

$$\text{CS1 : } f_{ij}(c_{ij} - p_{ij}\beta_j - \alpha_i) = 0 \quad \forall ij \quad (1)$$

$$\text{CS2 : } \alpha_i(a_i - \sum_j f_{ij}) = 0 \quad \forall i \quad (2)$$

$$\text{CS3 : } \beta_j(b_j - \sum_i p_{ij} f_{ij}) = 0 \quad \forall j \quad (3)$$

In each iteration, the algorithm 2.2 **Auction** picks a source, i , with positive surplus, $s_i = a_i - \sum_j f_{ij}$, and $\alpha_i > 0$. It then picks a sink j that provides it the largest effective profit, where the profit is specified by the quantity, $c_{ij} - p_{ij}\beta_j$. If the sink is not saturated flow is pushed from source i to sink j . If sink j is saturated, flow from a source i' , that is shipping flow to sink j at a lower of value β'_j , is replaced by flow from i , i.e. source i wins over the shipment to sink j from source i' by bidding a higher value (β_j) for j . After the change of flow assignment, it is checked (2.3 **UpdateBeta**(β_j)) that there exists at least one source shipping flow at the lower value. If not, then β_j is raised, indicating that shipping flow to j requires paying a higher value (and lower effective profit). At this stage, the process is repeated considering sinks in order of effective profit until

- (i) either all the surplus at node i is pushed out or
- (ii) $\alpha_i = 0$

Algorithm 2.1 Initialize

- 1: $\alpha_i \leftarrow \max_j c_{ij}$
 - 2: $f_{ij} = 0$
 - 3: $\beta'_j \leftarrow 0$
 - 4: $\beta_j \leftarrow 0$
-

To prove the correctness, we show that primal and dual equations are satisfied and complementary slackness conditions hold. All line numbers in the discussion below correspond to the description of **Algorithm 2.2 Auction**, except when stated.

Lemma 2.1 *The dual solution maintained by the algorithm is always feasible.*

Proof: We observe that the value of β_j never decreases as the algorithm progresses. Therefore, $\alpha_i < c_{ij} - p_{ij}\beta_j$ can only become true when α_i is modified. However, when α_i is modified, it is set to a value $\geq \max_j \{c_{ij} - p_{ij}\beta_j\}$ in line 24. At that point, $\alpha_i \geq c_{ij} - p_{ij}\beta_j$ for all j . \square
For use in the next lemma, we define the *total price of the incoming flow on sink j* to be $\sum_i p_{ij} f_{ij}$.

Lemma 2.2 *The primal solution maintained by the algorithm does not exceed the budget constraints.*

Algorithm 2.2 Auction

```
1: while there exists a source  $i$  such that  $\alpha_i > 0$  and  $\sum_j f_{ij} < a_i$  do
2:   Pick a sink  $j$  such that  $c_{ij} - p_{ij}\beta_j$  is maximized
3:   if the sink is saturated, i.e.  $\sum_i p_{ij}f_{ij} = b_j$  then
4:     Find a source  $i'$ , currently assigned at the lower value level i.e.  $y_{i'j} = \beta'_j$ 
5:     if  $i' \neq i$  then
6:       //determine the amount that can be replaced
7:        $x \leftarrow \min(s_i, f_{i'j}p_{i'j}/p_{ij})$ 
8:       // replace the flow
9:        $f_{ij} \leftarrow f_{ij} + x$ 
10:       $f_{i'j} \leftarrow f_{i'j} - xp_{ij}/p_{i'j}$ 
11:       $y_{ij} \leftarrow \beta_j$ 
12:    else
13:      //i is already shipping to j at the lowest level
14:      //Raise the assigned value without changing any flow
15:       $y_{ij} \leftarrow \beta_j$ 
16:    end if
17:    Update  $\beta_j$  and  $\beta'_j$ 
18:  else
19:    //The sink is unsaturated, push maximum flow possible
20:    //under supply and demand constraint
21:     $x \leftarrow \min(s_i, d_j/p_{ij})$ 
22:     $f_{ij} \leftarrow f_{ij} + x$ 
23:  end if
24:   $\alpha_i \leftarrow \max_j(c_{ij} - p_{ij}\beta_j, 0)$ 
25:  if  $\alpha_i = 0$  then
26:     $y_{ij} \leftarrow \beta'_j$  for all  $j$ 
27:    //This ensures there is no further rise in  $\beta_j$  without reducing  $f_{ij}$  to zero
28:  end if
29: end while
```

Algorithm 2.3 UpdateBeta (β_j)

```
1: if  $\beta_j = 0$  then
2:    $\beta_j \leftarrow \epsilon \min_i c_{ij}/p_{ij}$ 
3: else if  $\forall i : f_{ij} > 0, y_{ij} = \beta_j$  then
4:    $\beta'_j \leftarrow \beta_j$ 
5:    $\beta_j \leftarrow \beta_j(1 + \epsilon)$ 
6: else
7:   //no update required
8: end if
```

Proof: The flow on any edge is increased in lines 9 and 22. In line 9 the total price of the incoming flow on sink j increases by $x p_{ij}$. It is then reduced by $p_{i'j} x \frac{p_{ij}}{p_{i'j}}$ in line 10, and therefore, the budget is met.

In line 21, the total price increases by at most $d_j = b_j - \sum_i p_{ij} f_{ij}$ and thus the budget is not exceeded. \square

Lemma 2.3 *The primal solution maintained by the algorithm does not violate supply constraints.*

Proof: During the algorithm, the increase x in flow, f_{ij} , on an edge ij is always limited by the available supply (line 7 and 21) and therefore, the supply constraint is clearly satisfied. \square

Combining the above lemmas 2.2 and 2.3 we conclude:

Lemma 2.4 *The primal solution maintained by the algorithm is always feasible.*

Lemma 2.5 *During the course of the algorithm, if sink j is not saturated then $\beta_j = 0$. Further, $\forall j, \beta_j$ is non-decreasing during the course of the algorithm.*

Proof: The variable β_j for all j is initialized to zero. It is subsequently changed by the procedure **UpdateBeta**, which in turn is called only if the sink j is saturated. That is, line 17 is executed only when the “if” condition in line 3 : $\sum_i p_{ij} f_{ij} = b_j$ is satisfied. Moreover, $\forall j, \beta_j$ only increases in value. \square

The following lemma will be useful for bounding the complexity of the various algorithms:

Lemma 2.6 *Suppose f_{ij} , the flow on edge ij , that was pushed at a valuation of sink j , $y_{ij} = \beta_j / (1 + \epsilon)$, is reduced to zero, i.e. pushed back to zero. Then at the next occurrence of the event that corresponds to f_{ij} being pushed back on edge ij , the valuation of sink j , $y_{ij} \geq \beta_j$.*

Proof: This follows from the fact the β_j is monotonically non-decreasing. Flow is pushed back on an edge ij only when there is flow on the edge, f_{ij} that had been pushed at the companion valuation of $\beta_j / (1 + \epsilon)$, where β_j is the current valuation. Let flow be pushed back from sink j to source i on edge ij at some time t in the algorithm such that the flow is reduced to zero. Then when flow is pushed from i to j , at a subsequent time $t' > t$, the valuation $y_{ij} \geq \beta_j$ (line 11), since the valuation of j does not decrease. Consequently, if flow is pushed back again from j to i , the valuation has increased. \square

Lemma 2.7 *During the execution of the algorithm, the Complementary Slackness condition, CS1 is approximately satisfied for each edge ij . That is,*

$$\forall ij : f_{ij} > 0, \alpha_i \leq c_{ij} - p_{ij} \beta_j + \epsilon c_{ij}$$

Proof: When flow is pushed on the edge ij , $f_{ij} > 0$ and $\alpha_i = c_{ij} - p_{ij} \beta_j$ or is 0 (line 2 and 24 in **Algorithm 2.2 Auction**) and the condition is true. Note that at this event, $y_{ij} = \beta_j$. Subsequently, α_i or β_j may change. We consider the changes to α_i and β_j and to the flow, f_{ij} , between the event in the algorithm when positive flow is pushed on the edge ij and the event when this pushed flow reduces to zero, being pushed back on edge ij . In between the two events, since we assume that $f_{ij} > 0$, y_{ij} is either β_j or $\beta_j / (1 + \epsilon)$. We consider cases depending on whether $\alpha_i > 0$ or not.

1. $\alpha_i > 0$. Suppose α_i does not change but β_j has increased by a factor of $(1 + \epsilon)$. In this case y_{ij} has a value equal to β'_j , otherwise f_{ij} would be zero. Since the value of β_j changes by a factor of $(1 + \epsilon)$, $\alpha_i \leq c_{ij} - p_{ij}\beta_j + \epsilon c_{ij}$. Note that the value of β_j can not increase further unless f_{ij} reduces to zero.

If α_i has also changed, it only decreases, since β_j is increasing. Therefore, the inequality still holds.

2. $\alpha_i = 0$. No future change in the value of α_i can occur. Prior to the step at which α_i is set to zero, $\alpha_i = c_{ij} - p_{ij}\beta_j > 0$ for some j . This implies that $p_{ij}\beta_j < c_{ij}$. An increase in the value of β_j thus still results in $c_{ij} - p_{ij}\beta_j + \epsilon c_{ij} > 0$. Thus the inequality holds when α_i is set to zero. Further, when α_i becomes zero, the algorithm sets the value of y_{ij} to be β'_j so as to ensure that β_j is not increased without the event that flow is pushed back on ij and f_{ij} is reduced to zero. And, subsequently, source i is never considered for pushing flow f_{ij} since line 1 of **Algorithm 2.2 Auction** only considers vertices with $\alpha_i > 0$. Thus if flow $f_{ij} > 0$ then β_j has not changed. And if β_j has changed, the flow is set to zero. The inequality thus holds.

□

Furthermore,

Lemma 2.8 *Algorithm 2.2 Auction terminates.*

Proof: At each iteration of the algorithm, flow is pushed on an edge and $\forall j, \beta_j$ either remains the same or the value of β increases for some j (Lemma 2.5).

When flow is pushed on an edge ij , either (i) s_i goes to zero (ii) the budget at a sink d_j is met or (iii) the flow on a back edge ji' is reduced to zero. Note that, by lemma 2.6, in case (iii) flow cannot be pushed (in the forward direction) on the edge $i'j$ without increasing the value of y_{ij} .

The first event, case (i), results in flow being pushed from i to j such that $y_{ij} = \beta_j$ and will happen only once per source-sink pair until an event where flow is pushed back on an edge $j'i$ to add to surplus at i occurs. Thus this event occurs whenever new surplus is generated, which is bounded by the number of times flow is pushed back on an edge without case (iii) occurring. The number of times flow is pushed back on an edge, without reducing the flow to zero (case (iii)), is bounded since the amount of change of flow at every instance is a rational with bounded denominator.

The second event, case (ii), happens at most $O(m)$ times, since after that event the sink remains saturated during the subsequent operations of the algorithm. The third event, case (iii), occurs at most n times before all edges leading to a sink have flow reduced to zero, and consequently the value of β_j must rise. This bounds the total number of operations in the case that β_j remains the same. Further, the number of increases of β_j is bounded since β_j increases by a factor of $1 + \epsilon$ and its value is bounded above by $\max_j \{c_{ij}\}$. Thus the algorithm terminates in a finite number of steps. □

In order to show good convergence of the algorithm we introduce further modifications. However, when the algorithm terminates then the following is true:

Lemma 2.9 *At termination, the difference between the value of primal solution $\sum_{ij} c_{ij}f_{ij}$ and the value of the dual solution is at most $\epsilon \sum_{ij} c_{ij}f_{ij}$*

Proof: The value of the dual solution is

$$= \sum_i a_i \alpha_i + \sum_j b_j \beta_j \quad (4)$$

$$= \sum_i (a_i - \sum_j f_{ij}) \alpha_i + \sum_j (b_j - \sum_i p_{ij} f_{ij}) \beta_j + \sum_{ij} f_{ij} \alpha_i + \sum_{ij} f_{ij} p_{ij} \beta_j \quad (5)$$

When we subtract the value of the primal solution $\sum_{ij} c_{ij} f_{ij}$, from the above, the difference is

$$\sum_i (a_i - \sum_j f_{ij}) \alpha_i + \sum_j (b_j - \sum_i p_{ij} f_{ij}) \beta_j - \sum_{ij} f_{ij} (c_{ij} - \alpha_i - p_{ij} \beta_j) \quad (6)$$

Thus, the total absolute difference is at most $\Delta_1 + \Delta_2 + \Delta_3$ where

$$\Delta_1 = \left| \sum_{ij} f_{ij} (c_{ij} - \alpha_i - p_{ij} \beta_j) \right|$$

$$\Delta_2 = \left| \sum_i \alpha_i (a_i - \sum_j f_{ij}) \right|$$

$$\Delta_3 = \left| \sum_j \beta_j (b_j - \sum_i p_{ij} f_{ij}) \right|$$

From lemma 2.7, we have

$$\forall ij : f_{ij} > 0, \alpha_i \leq c_{ij} - p_{ij} \beta_j + \epsilon c_{ij}$$

Therefore,

$$\Delta_1 \leq \epsilon \sum_{ij} c_{ij} f_{ij}$$

From the termination condition of the algorithm, we know that for any source i such that $a_i - \sum_j f_{ij} > 0$, we have $\alpha_i = 0$.

Therefore

$$\Delta_2 = 0$$

For all unsaturated sinks $\beta_j = 0$. Therefore,

$$\Delta_3 = 0$$

Combining all of the above we get the lemma. \square

2.2 Modified Algorithm

In order to show the required complexity bound, and make its analysis simpler, we make a few modifications to the above algorithm. We address the details of the data structures required and present the complexity analysis. We also show that the modifications do not affect the correctness of algorithm 2.2 **Auction** as proved in section 2.

We introduce the concept of a preferred edge, back edge and derived graph as explained below.

Preferred Edge: For a given set of values for α_i and β_j we designate a preferred edge Pr_i for each source i . This edge maximizes $c_{ij} - p_{ij}\beta_j$, i.e. $Pr_i = ik$ where $k = \operatorname{argmax}_j (c_{ij} - p_{ij}\beta_j)$. In the case that multiple values of j satisfy the condition, only one is picked. The source i will be required to push flow along this edge until the edge is no longer is preferred.

Back-edge: All edges ij such that $f_{ij} > 0$ and $y_{ij} < \beta_j$. Flow is pushed back on these edges. The set of all such edges is termed B .

Derived Graph: We define the capacitated derived graph H with respect to an intermediate solution maintained by the modified algorithm. This graph consists of all the sources and sinks in the given bipartite graph and directed edges which are either the preferred edges or the back edges. All the preferred edges are oriented in the forward direction, whereas all the back-edges are oriented in the reverse direction, indicating the directions along which flow will be pushed. The capacity of a preferred edges is infinite whereas the capacity of a back-edge is the amount of flow carried on that particular edge. Formally, $H = (S, T, E, c)$ where

$$E = \{ij | i \in S, j \in T, ij \in Pr_i\} \cup \{ji, j \in T, i \in S, ij \in B\}$$

and the capacity function $c : E \rightarrow \mathbb{Z}^+$ is as follows:

$$c(ij) = \infty \quad \text{iff } ij \in Pr_i$$

$$c(ji) = f_{ij} \quad \text{iff } ij \in B$$

Before starting with the main algorithm, the following preprocessing step eases the description:

Pre-processing step: Consider a preferred edge $ij \in Pr_i$. If edge ij is also a back edge (thus creating a 2-cycle ($i \rightarrow j \rightarrow i$) in H) and is not the only back edge from j then the valuation at which flow is shipped from source i to sink j , i.e. $y_{ij} = \beta_j/(1 + \epsilon)$, is increased to β_j . Note that this modification is not done when the sink has only one back edge ij since the existence of at least one valuation equal to $\beta_j/(1 + \epsilon)$ does not allow a raise in the price of sink j . This ensures:

1. If there is a 2-cycle on sink j , sink j has only one back edge.
2. β_j does not change.

In the derived graph H , if $ij \in Pr_i$ before the modification then $ij \in Pr_i$ after the modification also, and source i will choose to push flow along its preferred edge.

The algorithm in section 2.1 pushed flow from a source to a sink causing flow to be pushed back to some other source and consequent independent processing of that other source was performed. The modified algorithm in this section, starting from a source with a surplus and positive α_i , finds a path along the edges of the derived graph H and pushes the flow according to the type of path discovered, thus ensuring that the surplus flow is processed in one push of flow along a flow path. The path found by the algorithm ends at either (i) an unsaturated sink, (ii) a source with $\alpha_i = 0$ or (iii) a cycle.

We detail the various cases that arise. A source i is picked such that $\alpha_i > 0$ and $\sum_j f_{ij} < a_i$. And a path, P , is discovered by, say a depth first search, along the edges of the derived graph H . There are different types of paths/cycles that may arise.




Figure 1: An example where path ends in a source with $\alpha_i = 0$




Figure 2: An example where path ends in a sink with $\beta_i = 0$




Figure 3: Example where the path ends in a cycle

1. **Type-1 paths:** We have two cases, one where the path finds a source i_k where $\alpha_i = 0$ and the second where the path finds an unsaturated sink j_k where $\beta_{j_k} = 0$. In the first case, let $P = i_0, j_0, \dots, i_k$ and in the second case let $P = i_0, j_0, \dots, j_k$. In both cases flow is pushed as follows: in the first case flow is pushed into source i_k (i.e. returned to the source) and in the second case flow is pushed into a sink (j_k), without violating feasibility or complementary slackness conditions. However, the flow may be limited by the capacity of the edges on the path, P . We define the capacity of the path, $c(P)$ as $c(P) = \min\{c(e) | e \in B \cap P\}$. The processing of the flow is as follows:
 - (a) The flow is limited by the surplus, s_{i_0} , at the starting source i_0 . Push all the surplus along the path. The surplus at the source disappears. This is true for both cases.
 - (b) The amount of flow is limited by the capacity of an unsaturated sink on path P (second case), i.e. $c(P) \geq d_{j_k}$. In this case, we push enough flow to saturate the sink and consequently, raise the value of β_j .
 - (c) The amount of flow is limited by the capacity of the path $c(P)$. In this case, we push flow from the source along the path. Let f be the flow starting at the source. Let $e = j_l i_{l+1}$ be the first edge such that $f > c(e) = c(j_l, i_{l+1})$. The amount of flow that can be pushed along this edge is limited by $c(e)$. Pushing flow along this edge reduces the capacity of the edge e to 0. The surplus $f - c(e)$ is added to the surplus of the previous source i_l which has $i_l j_l$ as the preferred edge. The flow pushed back through $j_l i_{l+1}$ is accumulated at i_{l+1} as a surplus. The surplus at i_{l+1} is pushed further along the path recursively. In the end one or more back edges would have zero flow and surplus may be left at each source occurring just prior to the sink where the back edge with zero flow is incident.
2. **Type-II paths:** The path ends at a 2-cycle (ji, ij) . A 2-cycle implies ji is the only back-edge on sink j , as soon as the flow is again pushed along edge ij , y_{ij} changes (line 15, **Algorithm 2.2 Auction**) and the only back edge disappears. This results in an increase in the value of β_j .

3. **Type-III paths:** A cycle is discovered since one of the sources is encountered again. Let the path discovered be denoted by $i_0, j_0 \dots i_k \dots j_q i_k$ which is decomposed into a simple cycle $C = i_k j_k \dots j_q i_k$ and a simple path $P_1 = i_0, j_0 \dots i_k$. The source i_k being the *entry point*, i.e., the first source that appears in the cycle. First the flow is pushed along the path $i_0 \dots i_k$ taking into account the capacities as in Type-I paths. Any additional surplus that appears at source i_k is handled as follows.

Consider an imaginary procedure that simulates sending flow around the cycle, C , multiple times, while conserving the flow at each of the intermediate sources in the cycle.

We define the *transfer ratio* ρ_l for a back-edge $j_l i_{l+1}$ w.r.t. C to be $p_{i_l j_l} / p_{i_{l+1} j_l}$ where $i_l j_l$ is the preferred edge for source i and $j_l i_{l+1}$ is a back edge.

Further define the *cumulative transfer ratio* ρ_l^c for the back-edge j_l, i_{l+1} with respect to a cycle C to be the amount that goes through the edge j_l, i_{l+1} if one unit of flow is pushed along the cycle starting at source i_k . This is equal to $\prod_{z=k}^l \rho_z$.

Let the cumulative transfer ratio of the cycle be $\rho_\odot = \prod_{l=k}^q \rho_l$. This is the amount of flow that reaches the first source, i_k , back after being pushed all around the cycle starting with a unit of flow at i_k .

Let us call the push of flow once around the cycle as one revolution. If we were to send the surplus at source i_k , s_{i_k} units of flow, repeatedly in this cycle for say, r times, while conserving the flow at each of the intermediate sources, the total flow shipped through the back edge $j_l i_{l+1}$ is given by the geometric series $F(l, r) = \sum_{z=0}^r s_{i_k} \rho_l^c (\rho_\odot)^z$. Since this back edge is capacitated, the maximum number of times flow can be pushed around the cycle before saturating this edge to capacity, which we term as the *limiting number of revolutions*, is given by $R_l = \arg\max_r (F(l, r) \leq f_{i_{l+1} j_l})$. This number can be computed in $O(1)$ time. It ignores that other edges may also limit the flow. We thus compute $R_{min} = \min_{l \in B \cap C} R_l$ to compute the *limiting number of revolutions for the cycle*, i.e. the smallest value of the limiting number of revolutions, computed from amongst the back edges in the cycle, C .

The flow on the edges of the cycle is modified as:

$$f_{i_l j_l} \leftarrow f_{i_l j_l} + \sum_{j=0}^{R_{min}} s_{i_1} \rho_{l-1}^c (\rho_\odot)^j$$

for forward edges $i_l j_l$ and

$$f_{i_{l+1} j_l} \leftarrow f_{i_{l+1} j_l} - \sum_{j=0}^{R_{min}} s_{i_1} \rho_l^c (\rho_\odot)^j$$

corresponding to back edges.

We consider cases based on the value of ρ_\odot :

- (a) $\rho_\odot < 1$. In this case we have a decreasing gain cycle, C , i.e. the value of the flow decreases as it is sent around the cycle. We compute the limiting number of revolutions for each back edge in the cycle, C . Note that this number may be infinity for some edges. However, this can be determined by computing the value to which the geometric series, $F(l,)$, converges and does not require simulating an infinite sequence of pushes.

After R_{min} rounds, the remaining surplus at i_k is reduced to $s_{i_k} \rho_{\odot}^{R_{min}}$. Note that if R_{min} is infinity then the surplus at the source, s_{i_k} decreases to zero. The surplus at every other source in the cycle remains unchanged.

Consider the cases depending on the value of the surplus at i_k :

- (i) the surplus remaining at i_k becomes zero, or,
 - (ii) the flow on one of the back edges reduces to zero. This case occurs when the number of revolutions is limited due to one of the back-edges. Since R_{min} integral revolutions need not achieve this, one more sequence of flow pushes around the cycle may be required.
- (b) $\rho_{\odot} \geq 1$. In this case the flow increases or remains the same as it goes along the cycle. There will be thus no decrease of surplus at the starting source vertex k . However, for each of the back-edges $R_l = \operatorname{argmax}_r (F(l, r) \leq f_{i_{l+1}j_l})$ is a finite number since each push of flow around the cycle reduces the capacity of the back-edges. The smallest of these values is chosen and the maximum number of rounds of flow to be pushed around the cycle is computed. Again, since R_{min} integral revolutions need not achieve this, it may require one more sequence of flow pushes around the cycle to reduce the flow on a back-edge to zero.

In the cases above where the surplus is not reduced to zero, at least one edge is saturated by the flow pushed around the cycle. Any surplus remaining at other nodes is pushed around the cycle to accumulate at the source(s) just before the saturated edge(s).

The overall structure of the algorithm is similar to that in section 2. After pushing the flow through the path or cycle discovered, values of β are updated if required, preferred and back edges computed as necessary and the algorithm repeats until all surplus is removed or until complementary slackness conditions are satisfied. This algorithm is termed as the *Modified Auction* algorithm.

Data Structures: In order to determine $\max_j c_{ij} - p_{ij}\beta_j$ for each i we maintain a heap. This heap can be updated when β_j changes in $O(\log m)$ time. The UpdateBeta procedure, as in the previous algorithm, utilizes a set data structure to determine the need for update of β_j in $O(1)$ time.

We next discuss how this approach improves the complexity. However, we can actually solve a more generalized problem, i.e. where each edge $(u, v), u \in S, v \in T$ has a capacity u_{ij} . So we defer a formal description and proofs of the above approach to section 3.

2.3 A discussion on the correctness and complexity of Modified Auction

We discuss the correctness and time complexity of the modified algorithm. Formal proofs are provided when we consider the generalized model including edge capacities.

Correctness: In order to justify that the modified algorithm terminates with a $(1 - \epsilon)$ approximate solution, first, we observe that while the algorithm 2.2 **Auction** does not require any particular order in which the surplus sources are picked, the sequence of choices of the sources from where the surplus is reduced, as imposed by the modifications, still ensures primal and dual constraints.

In the case that the algorithm encounters cycles, the modified algorithm essentially simulates multiple steps of the algorithm 2.2 **Auction**. The end result is exactly the same if each step

was performed individually and repeatedly. The amount of flow shipped out of a source at each step, however, may be *smaller* than what Algorithm 2.2 **Auction** would have shipped out. However, the primal and dual constraints are still ensured.

The modified algorithm makes the same changes to the variables α_i , β_j or f_{ij} as described in the previous section. The only modification is in the change to the variables, y_{ij} in the preprocessing procedure. This change does not affect the claim that $\beta_j \leq y_{ij}(1 + \epsilon)$ for all i in the proof of lemma 2.7 as y_{ij} is set to β_j itself. The pre-processing procedure leaves at least one back-edge intact; this ensures that no change in β_j takes place as a result of the preprocessing. All the arguments involving β_j are, therefore, unaffected. Thus primal and dual feasibility is maintained. Further the complementary slackness conditions are satisfied.

Complexity: For the complexity we consider the sequence of steps between two successive raises of the value of β , termed as a **phase**, in the algorithm.

We use the fact that for all j , the total number of times the value β_j rises is bounded. We show that each operation during the course of the algorithm can be charged to a rise in the value of β_j for some j . In fact, $\forall \beta_j$, every rise in the value of β_j is required to be charged at most $3n^2 + n$ times during the algorithm. We account for the work required for each push of flow along a path or cycle via a charging argument. There are at most n pushes of flow along edges of a path before either: (i) a 2-cycle is encountered, (ii) a source with $\alpha = 0$ is encountered, (iii) an unsaturated sink is encountered, or (iv) a cycle is encountered. Note that in the last case all the calculations for determining the limiting number of revolutions can be done in $O(n)$ time, since the cumulative transfer ratio can be determined by traversing the cycle once. As noted above, in all the cases that arise when flow is pushed around the cycle, either the surplus that has been pushed into the cycle reduces to 0 or the flow on one of the back edges is reduced to zero. The $O(n)$ amount of work can be charged to the disappearance of surplus at a source or reduction of flow to zero on a back-edge ji , and hence an increase in β_j , subsequently, when all back edges incident to the sink j have the corresponding flow through them reduced to zero. The other cases have a similar analysis. The detailed complexity of the algorithm will be described in the next section.

3 Edge Capacities

In this section we further generalize the Budgeted Transportation Problem. In this version each edge has an upper bound on the flow going through. The capacity of each edge is represented by $u_{ij} \in \mathbb{Z}^+$. We call this problem the Budgeted Transshipment (BTS) problem. It is easy to see that the budgeted transportation is a special case where edge capacities are infinite.

Consider the Linear Program for BTS

$$\text{maximize} \quad \sum_{ij} c_{ij} f_{ij} \tag{7}$$

subject to:

$$\sum_j f_{ij} \leq a_i \quad \forall i \quad (8)$$

$$\sum_i p_{ij} f_{ij} \leq b_j \quad \forall j \quad (9)$$

$$f_{ij} \leq u_{ij} \quad \forall ij \quad (10)$$

$$f_{ij} \geq 0 \quad \forall ij \quad (11)$$

The dual of the above program is

$$\text{minimize} \quad \sum_i \alpha_i a_i + \sum_j \beta_j b_j + \sum_{ij} u_{ij} \gamma_{ij} \quad (12)$$

subject to:

$$\alpha_i \geq c_{ij} - p_{ij} \beta_j - \gamma_{ij} \quad \forall ij \quad (13)$$

$$\alpha_i, \beta_j, \gamma_{ij} \geq 0 \quad \forall ij \quad (14)$$

And the complementary slackness conditions are

$$f_{ij}(c_{ij} - \alpha_i - p_{ij} \beta_j - \gamma_{ij}) = 0 \quad \forall ij \quad (15)$$

$$\alpha_i(a_i - \sum_j f_{ij}) = 0 \quad \forall i \quad (16)$$

$$\beta_j(b_j - \sum_i p_{ij} f_{ij}) = 0 \quad \forall j \quad (17)$$

$$\gamma_{ij}(u_{ij} - f_{ij}) = 0 \quad \forall ij \quad (18)$$

4 Approximation Algorithm for BTS

We now present a detailed approximation algorithm for the BTS problem.

We first (re)define some terms

Edge dual variable γ_{ij} : The dual variable associated with each edge. If ij is saturated, γ_{ij} is set as, $\gamma_{ij} = \min_{j'} \{(c_{ij} - p_{ij} \beta_j) - (c_{ij'} - p_{ij'} \beta_{j'})\}$ over all choices of edges ij' . It is set to 0 if edge ij is unsaturated.

Preferred Edge: For each source i an edge ij , termed Pr_i , such that ij is unsaturated and $c_{ij} - p_{ij} \beta_j$ is maximized. Note that in the case when more than one edge meet the criteria, one of them is chosen as the preferred edge. The preferred edges are directed forward (from i to j) in the derived graph. Therefore, every forward edge is unsaturated.

Back Edge: An edge ij such that $f_{ij} > 0$ and $y_{ij} = \frac{\beta_j}{1+\epsilon}$.

As in the previous algorithm, a derived graph H is used by the algorithm. The derived graph has vertices as the original vertex set and edges as back edges and preferred edges. The edges are directed with the preferred edge from source i , $Pr_i = ij$, being directed from source i to a sink j , with capacity given by $u_{ij} - f_{ij}$, and a back edge directed from a sink j to a source i' with capacity f_{ij} .

The algorithm differs from the one in section 2 in the use of the variables γ_{ij} . To ensure complementary slackness, this variable becomes non-zero only if edge ij is saturated. This variable is not explicitly maintained but its value can be deduced from α_i and β_j . Its value indicates the extra cost incurred in using another edge from source i . To ensure complementary slackness conditions, the value of the variable γ_{ij} implicitly reduces when β_j rises and when the reduced value becomes zero, the edge ij can be designated as a back-edge such that flow can be pushed back on the edge ij , leaving it unsaturated.

The following additions/modifications are made to the algorithm:

- During the initialization, γ_{ij} is set to zero.
- β_j update: This update occurs when $y_{ij} = \beta_j$, for all unsaturated edges with positive flow. See procedure 4.4 **update**.
- Preprocessing : Before the removal of 2-cycles, for each saturated edge ij , if $\gamma_{ij} > 0$, ij cannot be a back edge. In case $\gamma_{ij} \leq 0$, it becomes a back edge. (See procedure 4.4 **update**).
- Path/Cycle discovery and pushing of flow: In the BTS problem, the forward edges have a capacity as well. While pushing the flow in a path or in a cycle, the bottleneck edges could be forward edges in addition to back-edges. (See 4.7 **findpath**, 4.5 **pushFlowPath** and 4.6 **pushFlowCycle**).

The algorithm with the above modification is described as Algorithm 4.1 **Modified Auction**, along with the associated procedures.

The algorithm initializes the primal and dual solution using procedure 4.2 **Initialize**. Preprocessing is done in procedure 4.3 **Preprocess** to choose the preferred edge and eliminate 2-cycles.

It then repeats the following steps: find a path in the derived graph using 4.7 **findpath**; depending on whether the path found is a simple path or contains a cycle, the algorithm pushes flow using 4.5 **pushFlowPath** or 4.6 **pushFlowCycle**.

Procedure 4.5 **pushFlowPath** starts with the first source and moves flow along the path without accumulating it at any of the intermediate sources, unless, one of the edges on the path is saturated. In this case, the source preceding the edge accumulates some surplus.

Procedure 4.6 **pushFlowCycle** is more involved. It uses Procedure 4.5 **pushFlowPath** to transfer the surplus from the initial source to the first source that the algorithm encounters in the cyclic part of the path (See figure 4). It then computes the number of pushes of flow through the cycles that are required to (i) identify bottleneck edges or (ii) push the surplus to zero (See figure 5). Note that an extra last push may be required to transfer the flow from the first source in the cycle to the one just before the bottleneck edge as illustrated in figure 6.

When required, β_j is updated using Procedure 4.4 **update** and preprocessing is done for the next iteration. It also checks if a saturated edge ij which has a small value of γ_{ij} prior to the increase in β_j has the condition that $\gamma_{ij} = c_{ij} - p_{ij}\beta_j - \alpha_i < 0$, i.e. the edge is on the verge

(i.e. γ_{ij} is close to 0) of becoming unsaturated and allows for the edge to become unsaturated by designating the edge as a back-edge.

Each source maintains the effective profit $c_{ij} - p_{ij}\beta_j$ for all the unsaturated edges in a heap. Whenever an edge is saturated, it is removed from the heap. When γ_{ij} is set to 0, it is re-inserted into the heap. Each of these operations requires $O(\log m)$ time. These operations are performed when β_j changes, for any j .

The algorithm terminates when all the complementary slackness conditions are approximately met. The algorithm can be analyzed for correctness as well as complexity using the same framework as section 2.2.

The correctness is shown by proving that both primal and dual solutions are feasible throughout the execution and complementary slackness conditions are approximately met at termination.

The complexity is analyzed via a charging argument. Each iteration of the algorithm is charged to one of the events: clearing of surplus at a source, saturation of a sink or saturation of an edge. All such charges are then accounted for in the final analysis.

The details of the algorithm and the proofs are presented in the following section.

4.1 Detailed Algorithm

The modified algorithm for BTS is presented below in detail, followed by proof of correctness and complexity analysis. Since BTP is a special case of BTS, the algorithm is applicable to both the problems.

Algorithm 4.1 Modified Auction

```

1: Initialize
2: Preprocess
3: while there exists a source  $i$  such that  $\sum_j f_{ij} < a_i$  and  $\alpha_i > 0$  do
4:    $P \leftarrow \text{FindPath}(i)$ 
5:   if  $P$  contains a cycle  $C$  which is not a 2-cycle at the end of  $P$  then
6:     pushFlowPath ( $P'$ ) where  $P'$  is the portion of  $P$  before  $C$ 
7:     pushFlowCycle ( $C$ )
8:   else
9:     if  $P$  does not end with a 2-cycle then
10:      pushFlowPath ( $P$ )
11:    else
12:      //  $P$  ends at a 2-cycle  $i_l j_l$ 
13:      pushFlowPath ( $P'$ ) where  $P'$  excludes the last edge  $i_l j_l$ 
14:      Eliminate back edge  $j_l i_l$  and let  $y_{i_l j_l} \leftarrow \beta_{j_l}$ 
15:    end if
16:  end if
17:  Update  $\beta_j$ 
18: end while

```

4.2 Feasibility of the Solutions

Lemma 4.1 *The primal and dual solutions maintained by algorithm 4.1 Modified Auction are feasible throughout its execution.*

Algorithm 4.2 Initialize

- 1: $f_{ij} \leftarrow 0$ for all ij
 - 2: $\beta_j \leftarrow 0$ for all j
 - 3: $\gamma_{ij} \leftarrow 0$ for all ij
 - 4: $\alpha_i \leftarrow \max_j c_{ij}$ for all i
-

Algorithm 4.3 Preprocess

- 1: **for all** i such that $\sum_j f_{ij} < a_i$ **and** $\alpha_i > 0$ **do**
 - 2: Find an unsaturated ij such that $c_{ij} - p_{ij}\beta_j$ is maximized
 - 3: Make ij the preferred edge for source i
 - 4: $\alpha_i \leftarrow \max(0, c_{ij} - p_{ij}\beta_j)$
 - 5: //Remove all 2-cycles unless it involves the only back edge
 - 6: **if** $\exists i'$ such that $y_{i'j} < \beta_j$ **then**
 - 7: $y_{ij} \leftarrow \beta_j$
 - 8: **end if**
 - 9: **end for**
-

Algorithm 4.4 β_j update

- 1: **for all** j **do**
 - 2: //Update β_j if all assigned sources are shipping at the higher value
 - 3: **if** $\forall i$ such that edge ij is unsaturated and $y_{ij} = \beta_j$ **then**
 - 4: $\beta_j \leftarrow \beta_j(1 + \epsilon)$
 - 5: **for all** i such that edge ij is saturated **do**
 - 6: **if** $c_{ij} - p_{ij}\beta_j - \alpha_i < 0$ **then**
 - 7: //There is an edge with γ_{ij} close to zero, i.e. there is another edge with a larger effective profit
 - 8: Designate ij as a back-edge
 - 9: //Raise β_j to $\epsilon \min_i \frac{c_{ij}}{p_{ij}}$ if it is 0 and sink j is saturated
 - 10: **if** $\sum_i f_{ij}p_{ij} = b_j$ **and** $\beta_j = 0$ **then**
 - 11: $\beta_j \leftarrow \epsilon \min_i \frac{c_{ij}}{p_{ij}}$
 - 12: **end if**
 - 13: **end if**
 - 14: **end for**
 - 15: **end if**
 - 16: **Preprocess**
 - 17: **end for**
-

Algorithm 4.5 pushFlowPath(P)

Ensure: $P = \{i_1, j_1, i_2, j_2 \dots i_l, j_l\}$ or $P = \{i_1, j_1 \dots i_l, j_l, i_{l+1}\}$

- 1: ϕ is the amount that is transferred. Initially, it is the surplus at the first source in the path
 - 2: $\phi \leftarrow s_i$
 - 3: $r = l$
 - 4: **if** P ends at j_l **then**
 - 5: $r = l - 1$
 - 6: **end if**
 - 7: **for** $k \leftarrow 1$ to r **do**
 - 8: //The transferable amount is subject to the forward and backward edge capacity constraints
 - 9: $\phi \leftarrow \min(\phi, u_{i_k j_k} - f_{i_k j_k}, f_{i_{k+1} j_k} \frac{p_{i_{k+1} j_k}}{p_{i_k j_k}})$
 - 10: $f_{i_k j_k} \leftarrow f_{i_k j_k} + \phi$
 - 11: $y_{i_k j_k} \leftarrow \beta_{j_k}$
 - 12: // ϕ is multiplied by the fraction of two prices to reflect the proportional flow reduction on back edge $i_{k+1} j_k$
 - 13: $\phi \leftarrow \phi \frac{p_{i_k j_k}}{p_{i_{k+1} j_k}}$
 - 14: $f_{i_{k+1} j_k} \leftarrow f_{i_{k+1} j_k} - \phi$
 - 15: **end for**
 - 16: **if** $r = l - 1$ **then**
 - 17: $\phi \leftarrow \min(\phi, u_{i_l j_l} - f_{i_l j_l})$
 - 18: $f_{i_l j_l} \leftarrow f_{i_l j_l} + \phi$
 - 19: **end if**
-

Algorithm 4.6 pushFlowCycle(C)

Require: $C = \{i_1, j_1, i_2, j_2, \dots, i_k, j_k, i_{k+1} = i_1\}$ is a cycle in the derived graph

```

1: //Compute  $\rho_{\odot}$ 
2:  $\rho_{\odot} \leftarrow p_{i_k j_k} / p_{i_1 j_k} \prod_{z=1}^{k-1} \frac{p_{i_z j_z}}{p_{i_{z+1} j_z}}$ 
3: //Compute limiting number of revolutions for each edge
4: for  $l \leftarrow 1, k$  do
5:   if  $l > 1$  then
6:      $\rho_l'^c \leftarrow \prod_{z=1}^{l-1} \frac{p_{i_z j_z}}{p_{i_{z+1} j_z}}$ 
7:      $\rho_l^c \leftarrow \prod_{z=1}^l \frac{p_{i_z j_z}}{p_{i_{z+1} j_z}}$ 
8:   else
9:      $\rho_1'^c \leftarrow 1$ 
10:     $\rho_1^c \leftarrow p_{i_1 j_1} / p_{i_2 j_1}$ 
11:   end if
12:    $R_l' \leftarrow \operatorname{argmax}_r (\sum_{z=0}^r s_{i_1} \rho_l'^c (\rho_{\odot})^z \leq u_{ij} - f_{i_l j_l})$ 
13:    $R_l \leftarrow \operatorname{argmax}_r (\sum_{z=0}^r s_{i_1} \rho_l^c (\rho_{\odot})^z \leq f_{i_{l+1} j})$ 
14: end for
15:  $R_{min} \leftarrow \min_l (R_l, R_l')$ 
16: //  $R_{min}$  can be  $\infty$ 
17:  $l_{min}^1 \leftarrow \operatorname{argmin}_l (R_l')$ 
18:  $l_{min}^2 \leftarrow \operatorname{argmin}_l (R_l)$ 
19:  $l_{min} \leftarrow \min(l_{min}^1, l_{min}^2)$ 
20: //change the flow on each edge to simulate pushing of multiple revolutions
21: for  $z \leftarrow 1$  to  $k$  do
22:   if  $R_{min} \neq \infty$  then
23:      $f_{i_z j_z} \leftarrow f_{i_z j_z} + \sum_{j=0}^{R_{min}} s_{i_1} \rho_z'^c (\rho_{\odot})^j$ 
24:   else
25:     //Compute the convergence of the infinite series
26:      $f_{i_z j_z} \leftarrow f_{i_z j_z} + \sum_{j=0}^{\infty} s_{i_1} \rho_z'^c (\rho_{\odot})^j$ 
27:   end if
28:    $y_{i_z j_z} \leftarrow \beta_{j_z}$ 
29:    $f_{i_{z+1} j_z} \leftarrow f_{i_{z+1} j_z} - \sum_{j=0}^{R_{min}} s_{i_1} \rho_z^c (\rho_{\odot})^j$ 
30: end for
31: if  $R_{min}$  is finite then
32:   //Send one more cycle of flow ; this will saturate some edge in one revolution
33:   pushFlowPath( $i_1 \dots i_k j_k i_1$ )
34:   Let  $P \leftarrow \{i_1, j_1, \dots, i_{l-1}, j_{l-1}, i_l\}$  where the edge  $i_l j_l$  was saturated in previous step
35:   pushFlowPath( $P$ )
36: end if

```

Algorithm 4.7 findPath(i)

```

1: repeat
2:   if  $\alpha_i > 0$  or source  $i$  already belongs to  $P$  then
3:     return  $P$ 
4:   end if
5:   Add source  $i$  to  $P$ 
6:   Let  $j$  be the sink such that  $ij$  is Preferred( $i$ )
7:   if  $\beta_j = 0$  or sink  $j$  already belongs to  $P$  then
8:     return  $P$ 
9:   end if
10:  Add sink  $j$  to  $P$ 
11:  Let  $i'$  be a source such that  $i'j$  is a back edge
12:   $i \leftarrow i'$ 
13: until false

```




Figure 4: First we bring the surplus to the first source in the cycle, possibly saturating an edge en route

Proof: Consider the initial primal solution $f_{ij} = 0$ for all f_{ij} (line 1 in Procedure 4.2 **Initialize**). This is clearly a feasible primal solution.

Thereafter, the flow is increased or decreased in Procedures 4.5 **pushFlowPath** and 4.6 **pushFlowCycle**. Let us consider each possible change.

1. Increase/decrease of flow along a forward edge in line 9 of Procedure 4.5 **pushFlowPath**.
 - (a) The source constraint is satisfied as the increase in flow on edge $i_k j_k$ is at most the surplus generated at source i_k in the previous iteration.
 - (b) Also, the increase is at most the available edge capacity $u_{i_k j_k}$, therefore edge capacity constraint is satisfied.
 - (c) The total increase in the price of incoming flow on sink j is $p_{i_k j_k} \phi$ (line 10) and the total decrease is $p_{i_{k+1} j} \frac{p_{i_k j_k}}{p_{i_{k+1} j}} \phi$. Therefore, there is no net change and the budget constraint remains satisfied and tight.




Figure 5: Send the flow around the cycle as many times as possible without saturating any edge

- (d) The decrease on edge $i_k j_k$ is $\phi \frac{p_{i_k j_k}}{p_{i_{k+1} j_k}}$ (line 13, 19) where ϕ is at most $f_{i_{k+1} j_k} \frac{p_{i_{k+1} j_k}}{p_{i_k j_k}}$. Therefore the decrease is at most $f_{i_{k+1} j_k}$. Thus the positivity constraint on edge flows is satisfied.

2. Increase/Decrease in flow in Procedure 4.6 **pushFlowCycle**. Everytime the flow is pushed, the outgoing flow equals the incoming flow on every source and a proportional amount out of every sink. Therefore, not only are the source constraints satisfied, but also the budget constraints on the sink. The number of times the flow is pushed in the cycle is the minimum of the limiting number of revolutions. Therefore the capacity constraints are satisfied.

The last step uses Procedure 4.5 **pushFlowPath**, which satisfies the constraints as explained above.

Now consider the dual solution and its feasibility. For all i , α_i is initialized to $\max_j c_{ij}$ which is clearly larger than $\max_j c_{ij} - p_{ij}\beta_j - \gamma_{ij}$ as β_j and γ_{ij} are zero to begin with (lines 2 and 3 in Procedure 4.2 **Initialize**).

All the instances in the algorithm where these variables change are enumerated below. Assuming that the solution is feasible till these changes are made, we show the solution is still feasible after the change.

1. The update of β_j in Procedure 4.4 **β_j Update** at lines 4 and 11 increases it's value. Since any increase in β_j implies a decrease in the value $c_{ij} - p_{ij}\beta_j - \gamma_{ij}$, α_i remains larger than $c_{ij} - p_{ij}\beta_j - \gamma_{ij}$.
2. Change in γ_{ij} : γ_{ij} is assigned a value equal to $\max(0, \alpha_{ij} - c_{ij}' - p_{ij}\beta_j')$. Therefore, for every ij , $\alpha_{ij} \geq c_{ij} - p_{ij}\beta_j - \gamma_{ij}$ by construction.
3. Note that any change in α_i only occurs when there is a change in some β_j , it does not change independently. When it changes, it is set to $\max_{ij}(0, c_{ij} - p_{ij}\beta_j)$ (see line 4 in algorithm 4.3 **preprocess**). It is therefore, larger than $c_{ij} - p_{ij}\beta_j - \gamma_{ij}$ for all ij as $\gamma_{ij} \geq 0$.




Figure 6: Send it one more time possibly saturating an edge and accumulating surplus at the source before it

□

4.3 Complementary Slackness

We prove the following lemma about the satisfaction of complementary slackness conditions.

Lemma 4.2 *The algorithm terminates with the following conditions satisfied*

$$\forall \alpha_i > 0, \quad a_i - \sum_j f_{ij} = 0 \quad (19)$$

$$\forall \beta_j > 0, \quad b_j - \sum_i p_{ij} f_{ij} = 0 \quad (20)$$

$$\forall \gamma_{ij} > 0, \quad u_{ij} - f_{ij} = 0 \quad (21)$$

$$\forall f_{ij} > 0, \quad |c_{ij} - \alpha_i - p_{ij}\beta_j - \gamma_{ij}| \leq \epsilon c_{ij} \quad (22)$$

Proof: We consider each condition

1. Source slackness condition (19). These are the terminating condition for algorithm 4.1 (line 3). They are, therefore, satisfied at the end.
2. Sink slackness condition (20). Note that β_j is initialized to zero (Procedure 4.2, line 2). Thereafter, β_j is changed only when sink j is saturated (Procedure 4.4, line 11). Also, once a sink is saturated, it stays saturated. This can be seen by observing that any decrease of flow along any back-edge only takes place as a result of an increase in proportional amount of flow along some other edge on this sink.
3. Edge slackness condition (21). These are satisfied initially as $\gamma_{ij} = 0$. Thereafter, as long as γ_{ij} is greater than zero, the flow does not decrease on this edge as it is never designated

a back edge. Only when γ_{ij} is set to zero (Procedure 4.4, line 6), is the edge allowed to become a back edge.

4. Flow slackness condition (22). Similar to the proof of lemma 2.7, we show that the following is satisfied.

$$\forall f_{ij} > 0, \quad \alpha_i \leq c_{ij} - p_{ij}\beta_j - \gamma_{ij} + \epsilon c_{ij} \quad (23)$$

Consider unsaturated edges. Whenever a flow is increased along an edge ij it is a preferred edge. In Procedure 4.3, line 4, whenever an edge is chosen as a preferred edge, α_i is set to $c_{ij} - p_{ij}\beta_j$ with γ_{ij} being zero as it is an unsaturated edge. The slackness condition is therefore, satisfied. Subsequently, β_j can rise by a factor of $(1 + \epsilon)$ resulting in a rise of at most ϵc_{ij} in effective profit, still satisfying (23). Once β_j has changed, ij now becomes a back edge. Now, β_j can not change any further unless this back edge has zero flow (Procedure 4.4, lines 3,4 and 8). If as a consequence of a rise in β_j , α_i is set to zero, the inequality (23) is true and no further flow will be pushed from source i . β_j cannot rise until all flow is pushed back to source i .

Next consider when the edge is saturated. In this case γ_{ij} is set to $c_{ij} - p_{ij}\beta_j - \alpha_i$, thereby satisfying (23). However, flow may be pushed back on the saturated edge such that it becomes unsaturated and γ_{ij} is set to zero. This happens when β_j rises. The change in the quantity $c_{ij} - p_{ij}\beta_j$ is no more than ϵc_{ij} ; thus $|\gamma_{ij}| \leq \epsilon c_{ij}$ just prior to being set to zero. Thus condition (23) is satisfied after the change. Condition (22) immediately follows from condition (23).

Hence we conclude the above lemma. □

From the above two lemmas we conclude:

Theorem 4.1 *The algorithm determines a primal solution to the BTS problem such that $\sum_{ij} c_{ij} f_{ij} \geq (1 - \epsilon)OPT$ where OPT is the optimal solution to the BTS problem.*

Proof: The proof is similar to the proof of lemma 2.9. The value of the dual solution is

$$\begin{aligned} &= \sum_i a_i \alpha_i + \sum_j b_j \beta_j + \sum_{ij} u_{ij} \gamma_{ij} \quad (24) \\ &= \sum_i (a_i - \sum_j f_{ij}) \alpha_i + \sum_j (b_j - \sum_i p_{ij} f_{ij}) \beta_j + \sum_{ij} [(u_{ij} - f_{ij}) \gamma_{ij} + f_{ij} (\alpha_i + p_{ij} \beta_j + \gamma_{ij})] \quad (25) \end{aligned}$$

When we subtract the value of the primal solution $\sum_{ij} c_{ij} f_{ij}$, from the above, the difference is

$$\sum_i (a_i - \sum_j f_{ij}) \alpha_i + \sum_j (b_j - \sum_i p_{ij} f_{ij}) \beta_j - \sum_{ij} f_{ij} (c_{ij} - \alpha_i - p_{ij} \beta_j - \gamma_{ij}) \quad (26)$$

Thus, the total absolute difference is at most $\Delta_1 + \Delta_2 + \Delta_3 + \Delta_4$ where

$$\Delta_1 = \left| \sum_{ij} f_{ij} (c_{ij} - \alpha_i - p_{ij} \beta_j - \gamma_{ij}) \right|$$

$$\Delta_2 = \left| \sum_i \alpha_i (a_i - \sum_j f_{ij}) \right|$$

$$\Delta_3 = \left| \sum_j \beta_j (b_j - \sum_i p_{ij} f_{ij}) \right|$$

$$\Delta_4 = \left| \sum_{ij} \beta_j (u_{ij} - \sum_i \gamma_{ij} f_{ij}) \right|$$

By arguments similar to lemma 2.9 and using lemma 4.2 we get that the difference is at most $\epsilon c_{ij} f_{ij}$. □

4.4 Complexity

To prove the complexity we first consider the effect of the pre-processing step during the algorithm. In this step the algorithm removes every 2-cycle unless there is no other back-edge on the sink j .

The lemma below follows directly:

Lemma 4.3 *If there is a 2-cycle between a source i and a sink j then there is no back-edge incident onto j other than ij .*

The algorithm finds path and cycles and pushes flow on edges of the path or cycles. Each operation corresponding to a push of flow on an edge either enables a reduction in surplus and changes flow to the capacity of the edge. If the flow on a back edge reduces to zero, this leads to a change on the dual variables β_j for some sink j . We attempt to charge the various operation in finding paths or cycles and the subsequent processing to changes in the dual variables.

Lemma 4.4 *Each operation in the procedures **pushFlowPath**, **pushFlowCycle**, **findPath** can be charged to a rise in the value of β_j , for some j such that each rise in the value of β_j is charged $O(n^2)$ operations.*

Proof: We account for the work required for each push of flow along a path or cycle via a charging argument. Procedure **findPath** requires $O(n)$ steps. Further, during **pushFlowPath** there are at most n pushes of flow along edges of the path before one of these happens

1. A 2-cycle is encountered.

If the 2-cycle is between i and j , source i increases the assigned value y_{ij} to β_j . Since ji was the only back-edge on j (Lemma 4.3), there is a change in β_j . We charge the $O(n)$ pushes to this rise in β_j .

2. A source with $\alpha = 0$ is encountered. This case is analyzed together with the next case.
3. An unsaturated sink is encountered.

The previous cases can be considered together with this case as they are similar. In both these cases, the surplus from a node travels along the edges of the derived graph and reaches a source or a sink where no further push of flow is required. The following sub-cases arise:

- (a) The surplus at the starting node disappears and no new surplus is created

In this case, we charge the at most n pushes of flow along the edges of the path to the node at which the surplus disappears. We charge such a node only once for the

disappearance of the surplus. Note that any subsequent appearance of surplus will be charged to the creation of the surplus and is caused only in the case when the flow on a back-edge becomes zero, which is the second case explained below. Considering all possible source-sink pairs, there are, therefore, at most n^2 pushes charged to the nodes before either the surplus at each source is removed or a surplus appears.

- (b) There is a back edge ji such that f_{ij} becomes zero and a new surplus at i' is created where $i'j$ was the preferred edge for i' .

In this case we charge the push on edges of the path leading to this back edge, i.e. a charge of n , to the creation of the surplus at this back edge. By lemma 2.6 this back edge can not re-appear unless β_j changes. We allow an additional charge of n to pay for the possible disappearance of the flow from i' .

There are at most n back-edges on sink j . The value β_j changes when the flow on each of these edges is reduced to zero. Therefore, the total charge on all the back-edges incident out of a sink is at most $2n^2$ pushes of flow after which there is a change in the value β_j . This implies we have $2n^2$ pushes charged to a rise in β_j .

- (c) There is a forward edge ij which becomes saturated to capacity, i.e. $f_{ij} = u_{ij}$. If a forward edge is saturated, it stays saturated until it becomes a back edge when $\gamma_{ij} = 0$. When an edge ij becomes a back edge, by lemma 2.6 this edge cannot occur again as a back edge until β_j increases. We can, therefore charge this work to the corresponding rise in the value of β_j . The rise in β_j could be charged $O(n^2)$ times as in the previous case. Note that as a result of this saturating push, surplus could be generated on the source just before the edge ij . We put an additional charge of $O(n)$ on this edge to pay for the possible disappearance of this surplus in the future.

4. A cycle is encountered. In **pushFlowCycle**, all the calculations for determining the limiting number of revolutions can be done in $O(n)$ time, since the cumulative transfer ratio can be determined by traversing the cycle once. As noted above, in all the cases that arise when flow is pushed around the cycle, either (i) the surplus that has been pushed into the cycle reduces to 0 or (ii) the flow on one of the back edges is reduced to zero or (iii) the flow on one of the forward edges reaches capacity. The $O(n)$ amount of work can be charged to either the disappearance of surplus at a source, or a reduction of flow to zero on a back-edge ij , or the saturation of a forward edge. This implies an increase in β_j , subsequently, in the same way as described for simple paths above.

From the case analysis above, we conclude the total charge on a rise in β_j is $O(n)$ from case 1 and $O(n^2)$ from case 2,3 and 4, hence proving the lemma. \square

Lemma 4.5 *The algorithm terminates in $O(\epsilon^{-1}(n^2 + n \log m)m \log U)$ time, where $U = \max_{ij}(\frac{c_{ij}}{p_{ij}})/\epsilon \min_{ij}(\frac{c_{ij}}{p_{ij}})$.*

Proof: From lemma 4.4, there are $O(n^2)$ charges per rise in the value β_j . Once a phase is over, the change in value of β_j causes sources to update the heaps (one at each source). This takes no more than $O(n \log m)$ time.

The update procedure takes $O(1)$ amortized time by maintaining a set of values of y_{ij} , each value being either β_j or β'_j . Thus, we require $O(n \log m)$ at the end of a phase to update all the data structures. The preprocessing requires n steps.

Each rise causes the quantity β_j to grow by a factor of $1 + \epsilon$, starting with $\epsilon \min_i(\frac{c_{ij}}{p_{ij}})$. For each sink t_j such that $\beta_j > 0$, there exists an i such that s_i ships flow to t_j and thus

$0 \leq \alpha_i = c_{ij} - p_{ij}\beta_j$. This leads to $\beta_j \leq c_{ij}/p_{ij} \leq \max_{ij} \frac{c_{ij}}{p_{ij}}$. There are m different β_j . Therefore there can be no more than $O(m \log_{1+\epsilon} U)$ total changes in β_j for all j . Combining this with the fact that every rise in β_j is charged $O(n^2 + n \log m)$ amount of work and $\log_{1+\epsilon} U = O(\epsilon^{-1} \log U)$, we have the result. \square

5 Concave, Piecewise Linear Profit

We now describe how to extend the above algorithm to a profit function which is concave and piecewise linear. We use the common edge splitting technique in order to reduce the problem to the linear profit function. This transformation is very similar to the well known transformation of convex cost mincost flows to linear mincost flows, see [1]

Given an instance I of the problem with concave piecewise linear profits, we map it to an instance I' of the capacitated but linear profit version.

Let the profit function be defined as $c_{ijk} \in \mathbb{Z}$ for the edge ij and interval k each interval being of fixed length say l , total number of intervals being \mathcal{L} .




Figure 7: The profit function for an edge ij . The slope in an interval k is c_{ijk} . Number of intervals is \mathcal{L}

We then convert the problem by splitting each edge into \mathcal{L} edges, the k th edge being of capacity l and profit c_{ijk} . The price of each is the same as the one for edge ij .

It is easy to see that a solution to instance I' can be easily converted into a solution for instance I . A solution to I' , however, can not always be converted into a solution for I . Nevertheless, any solution to I' can always be modified into a solution with greater or same profit such that the new solution can be transformed into a solution for I .

If for a solution f_{ijk} the following property holds; $\nexists z_1, z_2$ such that $z_1 < z_2$ and $f_{ijz_1} < l$ and $f_{ijz_2} > 0$, then we can convert the solution to one for I .

If this property does not hold then the profit can be increased/kept the same while eliminating such pairs. Transfer $\min(l - f_{ijz_1}, f_{ijz_2})$ from the edge ijz_2 to ijz_1 . The profit $c_{ijz_1} \geq c_{ijz_2}$ because of concavity and thus we obtain a more profitable solution. Repeating these transfers, we get the property that $\nexists z_1, z_2$ such that $z_1 < z_2$ and $f_{ijz_1} < l$ and $f_{ijz_2} > 0$.

This approach can also be used to map an approximate solution. As a consequence, the algorithm 4.1 **Modified Auction** can be used to approximate the above problem after the ap-




Figure 8: Construction of multiple edges corresponding to one original edge.

appropriate transformation. The transformation, however, does involve an increase in the number of edges. We conclude with the following result.

Theorem 5.1 *The Budgeted Transportation problem where the profit function is piecewise-linear and concave, can be approximated within a factor of $(1-\epsilon)$ in $O(\epsilon^{-1}(n^2 + n \log m) \mathcal{L} m \log U)$ time where \mathcal{L} is the number of intervals in the profit function.*

6 Conclusion

We have presented an approximation scheme for Budgeted Transportation problem. We further generalize it to the capacitated version, which can then be used to solve the problem with a non-linear profit version.

The technique used is a kind of primal-dual mechanism based on Auctions. The variables are modified in small steps in order to maintain approximate slackness conditions. We use augmenting path based mechanism to improve the complexity of the scheme. However, because of the nature of the problem, the cycles may arise in addition to simple paths. We have shown that the cycles can be handled without too much work.

An interesting open question is whether it is possible to solve the problem exactly and as a consequence, generalized flow using the above scheme.

References

- [1] R.K. Ahuja, T.L. Magnanti, and J.B. Orlin. *Network Flows*. Prentice Hall, 1993.
- [2] N. Andelman and Y. Mansour. Auctions with budget constraints. In *SWAT*, pages 26–38, 2004.
- [3] D. Bertsekas. Auction algorithms for network flow problems: A tutorial introduction. *Computational Optimization and Applications*, 1:7–66, 1992.
- [4] D. Bertsekas, L. Polymenakos, and Paul Tseng. An ϵ -relaxation method for separable convex cost network flow problems. *SIAM J. Optim.*, 7(3):853–870, 1997.

- [5] D. P. Bertsekas and D. A. Castanon. The auction algorithm for the transportation problem. *Ann. Oper. Res.*, 20(1-4):67–96, 1989.
- [6] D. Bienstock and G. Iyengar. Approximating fractional packings and coverings in $o(1/\epsilon)$ iterations. *SIAM Journal on Computing*, 35(4):825–854, 2006.
- [7] L. Fleischer and K. Wayne. Fast and simple approximation schemes for generalized flow. In *SODA*, 1999.
- [8] N. Garg and J. Könemann. Faster and simpler algorithms for multicommodity flow and other fractional packing problems. In *IEEE Symposium on Foundations of Computer Science*, pages 300–309, 1998.
- [9] R. Garg and S. Kapoor. Auction algorithms for market equilibrium. In *STOC '04: Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 511–518. ACM Press, 2004.
- [10] R. Garg, S. Kapoor, and V. Vazirani. An auction based market equilibrium algorithm for the separable gross substitutability case. In *APPROX*, 2004.
- [11] A. Goldberg and R. Tarjan. Solving minimum-cost flow problems by successive approximation. In *STOC '87: Proceedings of the nineteenth annual ACM conference on Theory of computing*, pages 7–18. ACM Press, 1987.
- [12] A.V. Goldberg, S.A. Plotkin, and . Tardos. Combinatorial algorithms for the generalized circulation problem. *Math. Oper. Res.*, 16:351–379, 1991.
- [13] D. Goldfarb and Z. Jin. A faster combinatorial algorithm for the generalized circulation problem. *Math. Oper. Res.*, 21:529–539, 1996.
- [14] M. D. Grigoriadis and L. G. Khachiyan. Fast approximation schemes for convex programs with many blocks and coupling constraints. *SIAM Journal on Optimization*, 4:86–107, 1994.
- [15] F.L. Hitchcock. The distribution of a product from several sources to numerous localities. *Journal of Mathematical Physics*, 20(2):224–230, 1941.
- [16] Christos Koufogiannakis and Neal E. Young. Beating simplex for fractional packing and covering linear programs. In *FOCS '07: Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 494–504, Washington, DC, USA, 2007. IEEE Computer Society.
- [17] H.W. Kuhn. The hungarian method for the assignment problem. *Naval Research Quarterly*, 2(1 and 2):83–97, 1955.
- [18] Michael Luby and Noam Nisan. A parallel approximation algorithm for positive linear programming. In *STOC '93: Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 448–457, New York, NY, USA, 1993. ACM.
- [19] A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. Adwords and generalized online matching. In *FOCS*, 2005.
- [20] J.D. Oldham. Combinatorial approximation algorithms for generalized flow problems. In *SODA*, 1999.

- [21] C.H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Prentice Hall, 1982.
- [22] Serge A. Plotkin, David B. Shmoys, and Eva Tardos. Fast approximation algorithms for fractional packing and covering problems. *Math. Oper. Res.*, 20(2):257–301, 1995.
- [23] T. Radzik. Faster algorithms for the generalized network flow problem. *Math. Oper. Res.*, 23:69–100, 1998.
- [24] É. Tardos and K.D. Wayne. Simple generalized maximum flow algorithms. In *IPCO*, pages 310–324, 1998.
- [25] K. Wayne. A polynomial combinatorial algorithm for generalized minimum cost flow. In *STOC*, 1999.

7 Appendix-Relation to Generalized Flow

Apart from being a natural extension of the Transportation Problem BTP is also related to a well known flow problem, generalized flows. As mentioned above, it is a special case of Generalized Flow. It is interesting to see, however, that there also exists a reverse relationship. We show how to transform the generalized flow problem to Budgeted Transportation.

Mincost Generalized Flow: We are given a digraph $G(V, A)$, a cost function $c : A \rightarrow \mathbb{R}$, a capacity function $u : A \rightarrow \mathbb{R}^+$, a multiplier function $\mu : A \rightarrow \mathbb{R}^+$ a source $s \in V$, the supply at the source d_s , a sink $t \in V$ and the demand d_t at the sink. The goal is to find a flow function, $f : A \rightarrow \mathbb{R}^+$ such that the flow is conserved at the nodes, is multiplied across the arcs and meets the supply and demand constraints, while minimizing the total cost of the flow. This problem can be expressed as the following linear program.

$$\text{minimize} \quad \sum_{ij} c_{ij} f_{ij}$$

subject to :

$$\sum_i \mu_{ij} f_{ij} - \sum_k f_{jk} = 0 \quad \forall j \in V/\{s, t\} \quad (27)$$

$$\sum_k f_{sk} = d_s \quad (28)$$

$$\sum_i \mu_{it} f_{it} = d_t \quad (29)$$

$$f_{ij} \leq u_{ij} \quad \forall ij \in A \quad (30)$$

$$f_{ij} \geq 0 \quad \forall ij \in A \quad (31)$$

Mincost Budgeted Transportation: This problem is a more generalized version of the conventional transportation problem. We are given a bipartite graph B , consisting of sources $S = \{s_i\}$ and sinks $T = \{t_j\}$. A supply function $a : S \rightarrow \mathbb{R}^+$, a budget function $b_j \rightarrow \mathbb{R}^+$, a cost function $c : S \times T \rightarrow \mathbb{R}$ and a price function $p : S \times T \rightarrow \mathbb{R}$. The goal is to come up with

a flow function $f : S \times T \rightarrow \mathbb{R}^+$ such that the total flow going out of a source is equal to the supply, the total price of the flow coming into a sink is equal to the budget, and the cost of the flow is minimized. The problem is stated below as an LP.

$$\begin{aligned} & \text{minimize} && \sum_{i \in S, j \in T} c_{ij} f_{ij} \\ & \text{subject to :} \end{aligned}$$

$$\sum_{j \in T} f_{ij} = a_i \quad i \in S \quad (32)$$

$$\sum_{i \in S} p_{ij} f_{ij} = b_j \quad j \in T \quad (33)$$

$$f_{ij} \geq 0 \quad i \in S, j \in T \quad (34)$$

We first show that an MCGF problem instance transforms to a minimum Budgeted Transportation problem instance. The optima of the two instances are the same and the solutions themselves can be mapped to each other. The transformation is valid even for approximate solutions.

Let I be an instance of the MCGF problem. We transform it to an instance I' of Mincost Budgeted Transportation problem as follows

- For each node i in V , we have a corresponding source $s_i \in S$ with capacity $a_i = \sum_j u_{ij}$.
- For each arc (i, j) in A , we have a corresponding sink in T , t_{ij} , with budget $b_{ij} = u_{ij}$.
- We have one additional sink t_s with budget equal to the supply d_s in I . It is connected to the source corresponding to source node s with an edge whose price is 1 and cost is 0.
- Each sink t_{ij} has two incoming edges from s_i and s_j , and the cost and price are as follows

$$\begin{aligned} c'_{s_i, t_{ij}} &= 0 \\ p_{s_i, t_{ij}} &= 1 \\ c'_{s_j, t_{ij}} &= c_{ij} / \mu_{ij} \\ p_{s_j, t_{ij}} &= \frac{1}{\mu_{ij}} \end{aligned}$$

- The capacity of the source corresponding to the sink in I is equal to the demand d_t

Consider a pair of flow functions F and F' for the instances I and I' , respectively. Let these be related to each other according to the relations defined for all ij pairs in the instance I .

$$f'_{s_j, t_{ij}} = f_{ij} \mu_{ij} \quad (35)$$

$$f'_{s_j, t_{jk}} = u_{jk} - f_{jk} \quad (36)$$

Further, for the source s_s corresponding to the node s ,

$$f'_{s_s, t_s} = d_s$$

Lemma 7.1 *For a given solution F to the original instance I , F' is a feasible solution to I' of the same cost.*

Proof: We show that F' is a feasible solution to I' by showing that it meets the constraints 32 and 33.

For each source s_j , except the source s_s , the total outgoing flow is the sum of flows to the sinks corresponding to outgoing and incoming edges on node j , which is,

$$= \sum_i f'_{s_j, t_{ij}} + \sum_k f'_{s_j, t_{jk}}$$

Using the mapping of solutions, this is equal to

$$\begin{aligned} &= \sum_i f_{ij} \mu_{ij} + \sum_k (u_{jk} - f_{jk}) \\ &= \sum_k u_{jk} + \sum_i f_{ij} \mu_{ij} - \sum_k f_{jk} \\ &= \sum_k u_{jk} = a_j \quad [\text{using equation 27}] \end{aligned}$$

Therefore, constraint 32 is met.

For the source s_s , a similar analysis shows that constraint 32 is met, since $\sum_k f_{sk} = d_s$ and $f'_{s_s, t_s} = d_s$.

For each sink t_{ij} the total price of the incoming flow is

$$\begin{aligned} &f'_{s_i, t_{ij}} / \mu_{ij} + f'_{s_j, t_{ij}} \\ &= f_{ij} + u_{ij} - f_{ij} = u_{ij} \end{aligned}$$

Therefore, the constraints 33 are met.

Also, the total cost of F'

$$\sum_{t_{ij}} (0 \cdot f'_{s_i, t_{ij}} + c'_{s_j, t_{ij}} f'_{s_j, t_{ij}}) = \sum_{ij} f_{ij} \mu_{ij} c_{ij} / \mu_{ij} = \sum_{ij} c_{ij} f_{ij}$$

which is same as the cost of F . □

Lemma 7.2 *For a given solution F' to the instance I' , F is a feasible solution to I and is of the same cost.*

Proof: To show that the conservation of flow constraint is met, consider the total incoming flow on a node j , which is

$$\sum_i \mu_{ij} f_{ij} = \sum_i f'_{s_j, t_{ij}}$$

There are two sets of edges incident on source s_j . The ones corresponding to incoming edges on node j , (s_j, t_{ij}) and others that corresponding to outgoing edges (s_j, t_{jk}) . The sum of the flow on these is equal to a_j . So the the above is

$$= a_j - \sum_k f'_{s_j, t_{jk}}$$

Since the total price of incoming flow on sink t_{jk} is equal to $f'_{s_j, t_{jk}} + \frac{f'_{s_k, t_{jk}}}{\mu_{jk}}$ which is equal to the budget of the sink $t_{jk} = b_{jk}$, the above becomes,

$$\begin{aligned} &= a_j - \sum_k (b_{jk} - \frac{f'_{s_k, t_{jk}}}{\mu_{jk}}) \\ &= a_j - \sum_k u_{jk} + \sum_k f_{jk} \quad \text{since } b_{jk} = u_{jk} \\ &= \sum_k f_{jk} \quad \text{since } a_j = \sum_k u_{jk} \end{aligned}$$

which is equal to the outgoing flow $\sum_k f_{jk}$. For the case of the source and sink, there are no incoming or outgoing edges, respectively. But,

$$\sum_i \mu_{ij} f_{it} = \sum_i f'_{s_i, t_{it}} = d_t$$

thus ensuring the sink demand in I and

$$a_s - d_s = \sum_k f'_{s_s, t_{sk}} = \sum_k u_{sk} - \sum_k f_{sk}.$$

Thus the outgoing flow at the source,

$$\sum_k f_{sk} = d_s$$

since $a_s = \sum_k u_{sk}$.

Now we show that the capacity constraints are met. For each edge $(i, j) \in A$, from constraint 30, we have

$$\begin{aligned} f'_{s_i, t_{ij}} + \frac{f'_{s_j, t_{ij}}}{\mu_{ij}} &= b_{ij} = u_{ij} \\ \Rightarrow \frac{f'_{s_j, t_{ij}}}{\mu_{ij}} &\leq u_{ij} \quad \text{as } f'_{s_i, t_{ij}} \geq 0 \\ &\Rightarrow f_{ij} \leq u_{ij} \end{aligned}$$

The cost of the solutions are mapped exactly as in the proof of lemma 7.1 □

Bey combining the lemmas 7.1 and 7.2 we conclude the following theorem regarding the equivalence of the two problems.

Theorem 7.1 *Given an instance of the MCGF problem I we can construct an instance of Budgeted Transportation Problem such that the optimum solutions of each have the same cost.*

Also, note that the solutions themselves can be mapped to each other. Thus, by solving the Budgeted Transportation problem we not only determine value of the optimum solution, we can also construct the solution itself.

7.1 Mincost to Maxprofit Budgeted Transportation

In order to minimize the cost, we could maximize the negative of the cost $-c_{ij}$. The maximization, however, does not clear the sources and as such conservation of flow in the mapped solution to the original problem would not be achieved.

We can get over this problem by minimizing instead, $M - c_{ij}$ where M is a very large constant. Since every edge has a very large profit, the optimum will clear all the sources. Any error δ in clearance will produce a negative term $M\delta$ in the total profit. By choosing an M large enough, δ can be made negligibly small; i.e. smaller than the granularity of values in the solution of a Linear Program. We refer to [21] for the discussion and bound on this granularity.

The cost of the optimum would be $M \sum_i a_i - \sum_{ij} c_{ij} f_{ij}$, where $\sum_{ij} c_{ij} f_{ij}$ is an optimum solution for the mincost instance. We note that an approximate solution to Maxprofit Budgeted Transportation, would not immediately lead to an approximation algorithm for Generalized flow.